Naval Research Laboratory

Washington, DC 20375-5320

Algorithms for Improving the Overland Surveillance Capability of the E-2C AEW Aircraft

STEPHEN R. SCHUTZ
EDWARD M. KUTRZYBA

Radar Analysis Branch
Radar Division

December 8, 2000

20001220 163

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	December 8, 2000	June 1999 — October 2000

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
Algorithms for Improving the Overland Surveillance Capability of the E-2C AEW Aircraft	

6. AUTHOR(S)

Stephen R. Schutz and Edward M. Kutrzyba

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Naval Research Laboratory Washington, DC 20375-5320	NRL/MR/5315--00-8518

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER
Office of Naval Research 800 North Quincy Street Arlington, VA 22217	

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT	12b. DISTRIBUTION CODE
Approved for public release; distribution is unlimited.	

13. ABSTRACT (Maximum 200 words)

This report discusses two algorithms developed at the Naval Research Laboratory (NRL) with the goal of improving the overland littoral surveillance capability of airborne UHF radars such as the APS-145. The APS-145 is the surveillance radar deployed on the E-2C carrier based airborne early warning aircraft. The impetus for this study is the fact that radars generally have much greater accuracy in the downrange direction than in the crossrange direction and, in the case of the APS-145 radar, the downrange measurement uncertainty is small compared to the typical spacing between roads while the crossrange measurement uncertainty is large compared to the typical spacing between roads. The first algorithm employs Digitized Features Analysis Data (DFAD) to leverage the downrange accuracy of the radar to improve on-road target crossrange accuracy. The second algorithm takes advantage of the fact that the first algorithm 1) associates on-road targets with a particular road and 2) reduces the tracking problem from two dimensions to one. Both algorithms were developed and evaluated using simulated data. The algorithm for improving target location accuracy was also exercised using measured data.

14. SUBJECT TERMS		15. NUMBER OF PAGES
Overland littoral surveillance Multitarget tracking Ground target tracking Digitized features analysis data		18
		16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

CONTENTS

ALGORITHMS FOR IMPROVING THE OVERLAND SURVEILLANCE CAPABILITY OF THE E-2C AEW AIRCRAFT

INTRODUCTION

This report discusses two algorithms developed at the Naval Research Laboratory with the goal of improving the overland littoral surveillance capability of airborne UHF radars such as the APS-145.[1] The APS-145 is the surveillance radar deployed on the E-2C carrier based airborne early warning aircraft. The impetus for this study is the fact that radars generally have much greater accuracy in the downrange direction than in the crossrange direction and, in the case of the APS-145 radar, the downrange measurement uncertainty is small compared to the typical spacing between roads while the crossrange measurement uncertainty is large compared to the typical spacing between roads. The first algorithm employs Digitized Features Analysis Data (DFAD) to leverage the downrange accuracy of the radar to improve on-road target crossrange accuracy. The second algorithm takes advantage of the fact that the first algorithm 1) associates on-road targets with a particular road and 2) reduces the tracking problem from two dimensions to one. Both algorithms were developed and evaluated using simulated data. The algorithm for improving target location accuracy was also exercised using measured data.

TARGET LOCATION ACCURACY

Radars generally have much greater accuracy in the downrange direction than in the crossrange direction and, in the case of the APS-145 radar, the downrange measurement uncertainty is small compared to the typical spacing between roads while the crossrange measurement uncertainty is large compared to the typical spacing between roads. This report discusses a technique which has been developed and assessed, which utilizes DFAD road data to leverage the relatively high downrange accuracy of the radar to improve overall target location accuracy. The technique was evaluated via computer simulation. The computer simulation generates simulated DFAD road data points, computes the coordinates of on-road targets moving along the simulated roads with constant speed, and generates simulated detection coordinates by adding independent Gaussian noise samples, with different standard deviations in the range and bearing directions, to the actual target coordinates. An example of a simulated road network, target coordinates and detection coordinates is shown in Figure 1. The standard deviations of the Gaussian downrange and crossrange noise are not representative of any particular radar but were chosen for illustrative purposes.

The proposed algorithm determines the Maximum Likelihood Estimate (MLE) target coordinates subject to the constraint that the estimated target coordinates lie on a road. The likelihood function is of the form:

$$p(\hat{\vec{r}} \mid \vec{r}) = \alpha \, exp\left\{-\frac{1}{2}(\hat{\vec{r}} - \vec{r})^T \bar{R}^{-1}(\hat{\vec{r}} - \vec{r})\right\} = \alpha \, exp\left\{-\frac{1}{2}\left(\frac{\Delta^2_{Range}}{(StdRange)^2} + \frac{\Delta^2_{Bearing}}{(StdBearing)^2}\right)\right\} \quad (1)$$

Fig. 1 - Simulated road network, target coordinates, and detection coordinates

where: $p(\hat{r} \mid \vec{r})$ is the probability that the actual target coordinates are $\hat{r}$ given that the detection coordinates are $\vec{r}$, $\bar{R}$ is the covariance matrix, $\Delta_{Range\ (Bearing)}$ is the component of $(\hat{r} - \vec{r})$ in the range (bearing) direction, and *StdRange (StdBearing)* is the range (bearing) accuracy. The superscript ^{T} denotes transpose. The MLE of $\hat{r}$ is the estimated position that maximizes $p(\hat{r} \mid \vec{r})$. Defining normalized Range-Bearing coordinates as $\dfrac{Range}{StdRange}$ and $\dfrac{Bearing}{StdBearing}$, the MLE of $\hat{r}$ corresponds to the point on a road closest to the detection coordinates, measured in normalized Range-Bearing coordinates.

Illustrated in Figure 2 is the algorithm by which the MLE target position on a road is found. Plotted in the Figure are a detection (diamond) and several consecutive DFAD road points (squares). Positions are plotted in normalized Range-Bearing coordinates. The stretches of road between consecutive DFAD points are modeled as straight lines and are referred to as road segments. Since it is not known a priori on which road segment the MLE target position will lie, the MLE target position for every road segment is found. The filled circles are the MLE target positions on two of the DFAD road segments. Two cases need to be considered. Case 1 is illustrated by the triangle formed by the detection and the rightmost road segment. The MLE target position is indicated by the filled circle along the base of the triangle. For Case 1 the MLE target position, the point on the road segment closest (measured in normalized

Range-Bearing coordinates) to the detection coordinates, is at the foot of the altitude of the triangle. Case 2 is illustrated by the triangle formed by the same detection and the adjacent road segment. For Case 2 the foot of the altitude lies outside the triangle and the MLE target position is one or the other of the two DFAD road points.
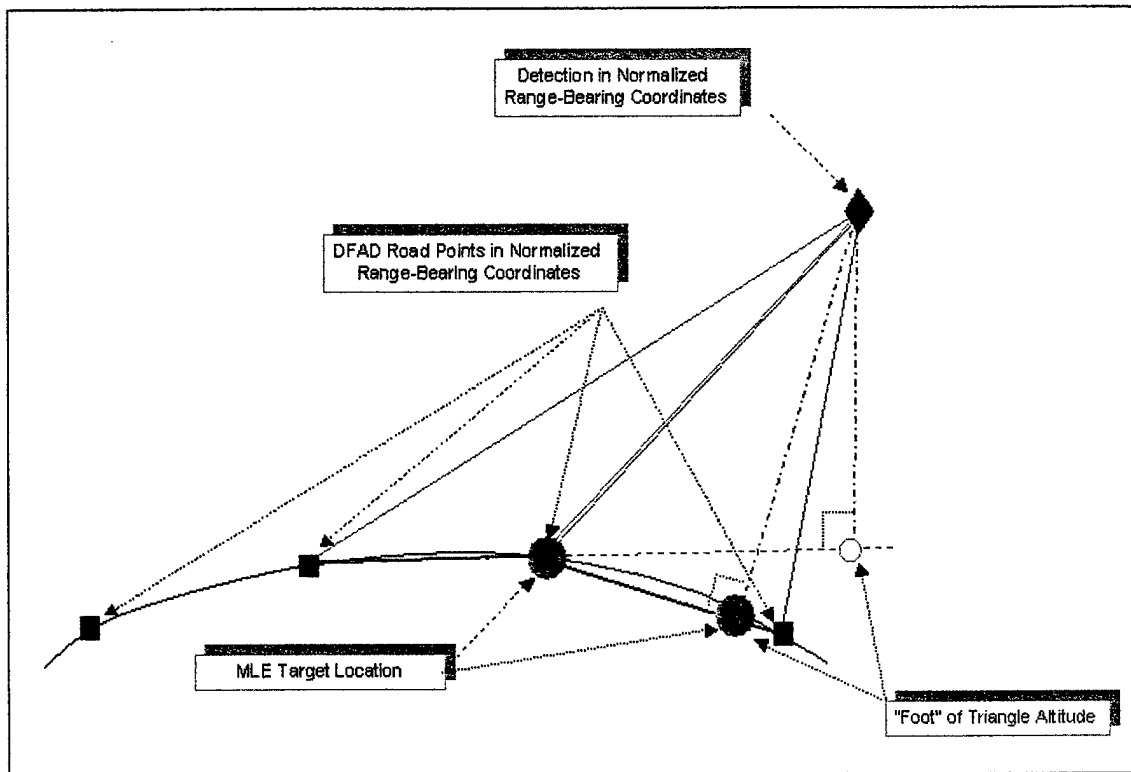

Fig. 2 - Illustration of MLE algorithm

Illustrated in Figure 3 is the test used to determine whether the geometry of a particular detection and a particular road segment corresponds to Case 1 or Case 2. If the geometry corresponds to Case 1, the altitude $\overrightarrow{AD}$ will lie within triangle $\overline{ABC}$ and the dot product $\overrightarrow{AB} \bullet \overrightarrow{BC}$ and the dot product $\overrightarrow{AC} \bullet \overrightarrow{BC}$ will have opposite signs. For this case the MLE position is point D at the foot of altitude $\overrightarrow{AD}$. The coordinates of point D are found by first noting that the area of triangle $\overline{ABC}$ is given by one-half the magnitude of the cross product of any two sides of the triangle, and also by one-half the base times the altitude, as indicated in Figure 3. The coordinates of point D are given by the coordinates of point B plus vector $\overrightarrow{BD}$. If the dot products $\overrightarrow{AB} \bullet \overrightarrow{BC}$ and $\overrightarrow{AC} \bullet \overrightarrow{BC}$ have the same signs, then the MLE point on that road segment is the DFAD road point closest to the detection in normalized Range-Bearing coordinates.

The MLE target positions for each detection for every road segment on every road is determined in normalized Range-Bearing coordinates. A plot showing the normalized distances from a particular detection to the closest points on all of the road segments on each of the roads in the road network is shown in Figure 4. The results shown in the Figure are for simulated data. The normalized distances are proportional to the negative of the logarithm of the likelihood function and so the smaller the normalized

Area of $\triangle ABC = \frac{1}{2}\left|\vec{AB} \times \vec{AC}\right| = \frac{1}{2}\left|\vec{BC}\right|\left|\vec{AD}\right|$

$$\left|\vec{AD}\right| = \frac{\left|\overline{AB} \times \overline{AC}\right|}{\left|\overline{BC}\right|}$$

$$\left|\vec{BD}\right| = \sqrt{\left|\vec{AB}\right|^2 - \left|\vec{AD}\right|^2}$$

$$\vec{BD} = \left(\frac{\sqrt{\left|\vec{AB}\right|^2 - \left|\vec{AD}\right|^2}}{\left|\vec{BC}\right|}\right)\vec{BC}$$

Altitude $\overline{AD}$ Lies Within $\triangle ABC$ If :

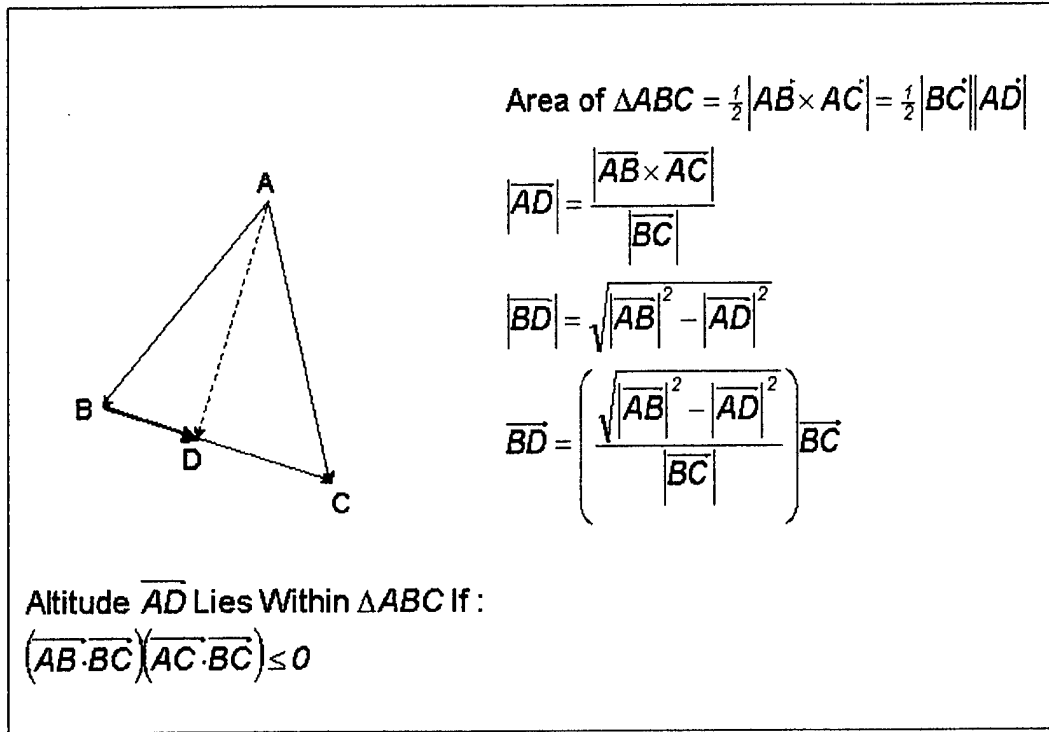$$\left(\overline{AB}\cdot\overline{BC}\right)\left(\overline{AC}\cdot\overline{BC}\right) \leq 0$$

Fig. 3 - Test for case 1 or case 2

distance, the higher the probability that the target lies at that position. As one would expect, the plot exhibits local minima which correspond to the MLE target positions on each road. The global minimum corresponds to the MLE target position for the entire road network. The global and local minimum normalized distances can be used to decide whether a particular detection is off-road or on-road and if on-road, which road. It is useful to determine the MLE target positions on each road in the event that collateral information indicates that the target is on a particular road.

Shown in Figure 5 is a simulated road network consisting of three roads, a detection (diamond), and the MLE target positions on each of the three roads, in normalized Range-Bearing coordinates. It can be seen that the MLE positions on roads 1 and 3 correspond to Case 1 (altitude lies within triangle), and that the MLE position on road 2 corresponds to Case 2 (the MLE position coincides with a DFAD road point). The insert at the upper left shows what the road network, detection, and MLE points look like when plotted in equally scaled rectangular coordinates.

Plotted in Figure 6 is the road network, the ground truth for five on-road targets over five scans (triangles), the corresponding detection coordinates (diamonds), and the global MLE target positions (asterisks). It can be seen that the algorithm correctly maps MLE points identified in normalized Range-Bearing coordinates onto positions on the road network. It can also be seen that the algorithm greatly improves the target location accuracy for the majority of the detections. However, the target location accuracy is degraded for a significant number of the detections. The degradation occurs in those cases where the global MLE target position lies on the wrong road.

Fig. 4 - Normalized distances from one detection to all road segments

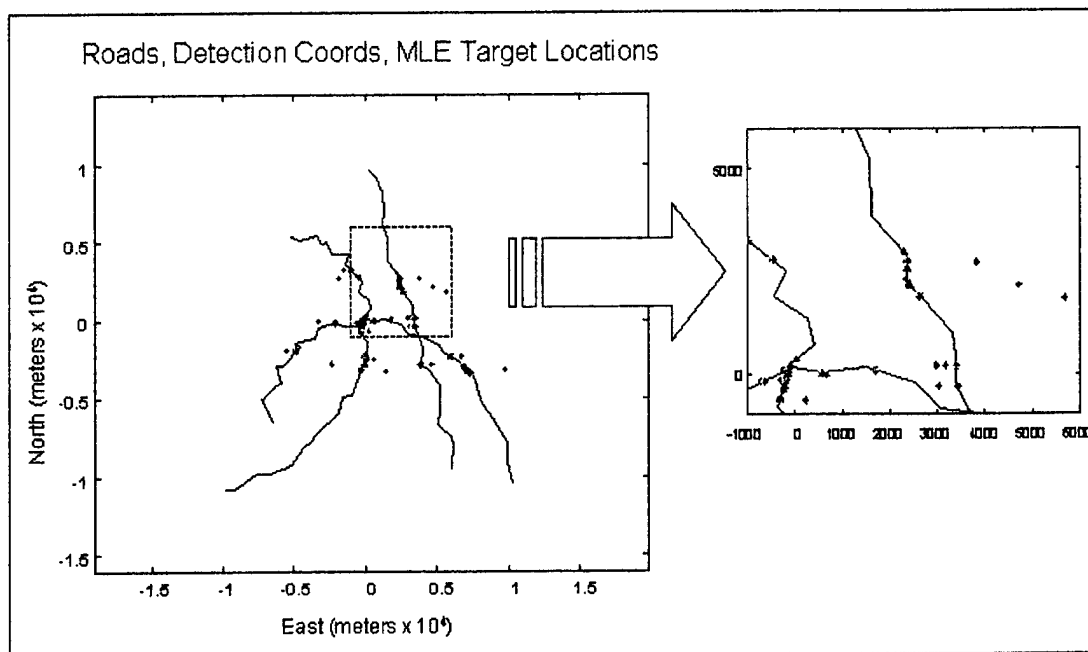Fig. 5 - Local MLE target positions in normalized range-bearing coordinates

Fig 6 - DFAD road network, detections, and global MLE positions in equally scaled rectangular coordinates

The nature of the problem is illustrated in Figure 7. Shown in the Figure are the simulated ground truth target positions for five targets, the coordinates of five detections, and the MLE target positions. Note that the target is moving along a North-South road. The platform is to the south and the uncertainty in the bearing measurement is manifested in East-West scatter. For two of the detections, the global MLE target positions lie on the correct road and target location accuracy is greatly improved. For the three other detections, the global MLE target positions lie on the wrong road. From the Figure it is clear that if the uncertainty in the bearing measurement exceeds the separation between roads, the global MLE target position is likely to be on the wrong road. There are several approaches that can be taken to remedy this type of error. First, rather than considering only the global MLE target positions, several local MLE positions for each detection could be considered. If collateral information (e.g. subsequent scan detection positions, INTEL, etc.) indicates that a target is on a particular road, then rather than using the global MLE position, the local MLE position for that target on that road could be used instead. Second, Minimum Detectable Velocity (MDV) considerations can be used to discount MLE positions on certain roads.

If the orientation of a road segment is such that the component of target velocity along the line-of-sight (LOS) is smaller than the MDV, no detections should be associated with that road segment. The LOS component of target velocity is a function of the unknown target speed and of the known road segment-LOS geometry. Using the known information, the algorithm computes the component of the road segment direction along the LOS and discounts MLE target positions on road segments whose LOS components are below some threshold. Comparison of Figure 8 with Figure 7 illustrates the effect of increasing that threshold. In Figure 7, two of the global MLE positions lie on a road that runs primarily perpendicular to the LOS. The results shown in Figure 7 were generated using a low LOS component

Fig. 7 - DFAD road points, detections, and MLE target positions with low MDV threshold

threshold. The results shown in Figure 8 were generated using a large LOS component threshold. It can be seen that the two global MLE positions along the East-West road for two of the detections in Figure 7 are replaced by MLE positions along North-South roads. One of these new MLE positions is now on the correct road and is a much more accurate estimate of target position while the other new MLE position still lies on the wrong road and does represent an improvement in target location accuracy. On balance, it would appear that the use of MDV considerations could be helpful in improving the target location accuracy of the technique.

The results discussed above were obtained using simulated data. Results obtained using measured APS-145 data are discussed below. The available APS-145 data did not come with collateral surface target ground truth data and so no assessment of target location accuracy could be made. However, the measured APS-145 data were valuable for purposes of algorithm development. One scan of processed APS-145 data with detections indicated by diamonds is shown in Figure 9. The data were recorded in the vicinity of Sedalia Missouri at 3:16 PM on May 20, 1998. Data over a large sector in the left half of the scan were corrupted by propeller modulation and discarded. The DFAD road network is also plotted in the Figure. In this scan there were 314 radar detections in the Doppler bins corresponding to LOS speeds greater than 35 statute mph. The range to the center of the patch was 78.1 nmi. The platform was nearly due South of the patch.

Fig. 8 - DFAD road points, detections, and MLE target positions with high MDV threshold

Fig. 9 - DFAD road data overlaid on measured APS-145 data

A cell and sector structure was developed to reduce the number of road segments processed in conjunction with each detection. In Range-Bearing coordinates a cell is a rectangle with dimensions on the order of a few standard deviations in range and bearing. A sector consists of a cell and the eight surrounding cells. If the coordinates of a detection lie within a cell, the actual target position is, with high probability, within the corresponding sector. The algorithm is applied cell by cell, sector by sector such that for detections in a particular cell, only DFAD road points within the corresponding sector are considered. The compartmentalization afforded by the cell and sector structure greatly speeds up processing. However, by considering only DFAD road points within a sector, road segments that extend to a DFAD point outside the sector, are not considered. Under certain circumstances this could result in the MLE target position lying on the wrong road with the concomitant degradation in accuracy.

Shown in Figure 10 are the results for a portion of the APS-145 data. The cells appear as sectors of annuli in Latitude-Longitude coordinates. The detection coordinates are indicated by diamonds and the global MLE target positions by asterisks. The axes of this plot are Latitude and Longitude. As a result, the cells and sectors which are rectangular in Range-Bearing coordinates appear as sectors of annuli in this plot. These results indicate that the technique functions properly in the sense that the MLE target positions do indeed lie on DFAD roads. Since ground truth data are not available, little can be said, one way or the other, regarding the performance of the technique with regard to accuracy of the MLE target locations.


Fig. 10 - DFAD road data, APS-145 detections, and MLE target positions

The potential improvement in E-2C/APS-145 ground target location accuracy achievable through implementation of this MLE technique was estimated via computer simulation. The technique was seen to work very well although it is susceptible to large errors when the global MLE target position lies on the wrong road. This type of error is essentially unavoidable if 1) the uncertainty in radar bearing measurement is large compared to the separation between the roads, or 2) no collateral information is available to associate a target to a particular road. It is of interest to estimate the improvement in target location in the absence of this type of error. The optimal geometry has the road parallel to the LOS so that a large component of target Doppler is seen, and the scatter in the data is essentially perpendicular to the road. The worst case has the road perpendicular to the line of sight. For this case the technique provides essentially no improvement in target location accuracy (MDV considerations make this moot). An analysis was performed using values for azimuthal and downrange accuracy appropriate to the APS-145, for a variety of road orientations. This analysis suggests that so long as the target is associated with the correct road, the target location accuracy of the APS-145 can be significantly improved.

TRACK ASSOCIATION

A second algorithm, a track association algorithm, was developed for use in conjunction with the target location algorithm discussed above. This algorithm could be used by the E-2C/APS-145 for tracking on-road targets over multiple scans, in a high target density environment. The track association algorithm exploits two important features of the target location algorithm. 1) The target location algorithm uses DFAD road data to convert a two-dimensional problem in Range-Bearing space into a one-dimensional problem with the dimension being displacement along a road. 2) The target location algorithm partitions the detections among the roads, converting a single problem with a large number of detections into several independent problems each with a smaller number of detections to deal with.

The underlying assumption on which the track association algorithm is based is that targets travel along roads with constant speed. If the time interval between scans is constant, then the change in displacement along a road, for such a target, should be the same from scan to scan. The algorithm examines all possible scan-to-scan associations of estimated target positions or tracks, calculates the standard deviations of the scan-to-scan displacements for each such track, and sorts those tracks in the order of increasing standard deviation. It is assumed that the smaller the standard deviation, the higher the probability that the track is true.

The algorithm selects the most probable tracks by first generating a track table whose columns represent all potential tracks. The track table is generated row by row as illustrated in Figure 11. Each entry in the table is the displacement along a road of an MLE target position. In the Figure the table entries are denoted by a letter and a number. The number refers to the scan in which the detection was made. The letter is an arbitrary designation used for labeling the detections in a scan. Thus the m^{th} row of the table contains only the displacements measured in the m^{th} scan. Initially, the first row consists of the displacements of the targets detected along the road under consideration, for the first scan. Next, the first row is replaced by the Kronecker product of the (initial) first row times a row vector comprised of N_2 ones, where N_j is the number of detections in scan number j. A second row, given by the Kronecker product of N_1 ones times the displacements of targets detected along the road in scan number 2 is appended to the first row. Thus the first row is stretched out, with each element repeated N_2 times, while the second row is expanded by replication. Next, the first two rows are replaced by the Kronecker product of the (current) first two rows times a row vector consisting of N_3 ones while the third row is formed by the Kronecker product of $(N_1 \times N_2)$ ones times the displacements of targets detected in scan number 3. The procedure is repeated until the track table contains M rows where M is the number of scans. Every possible track is represented by a column in the table.
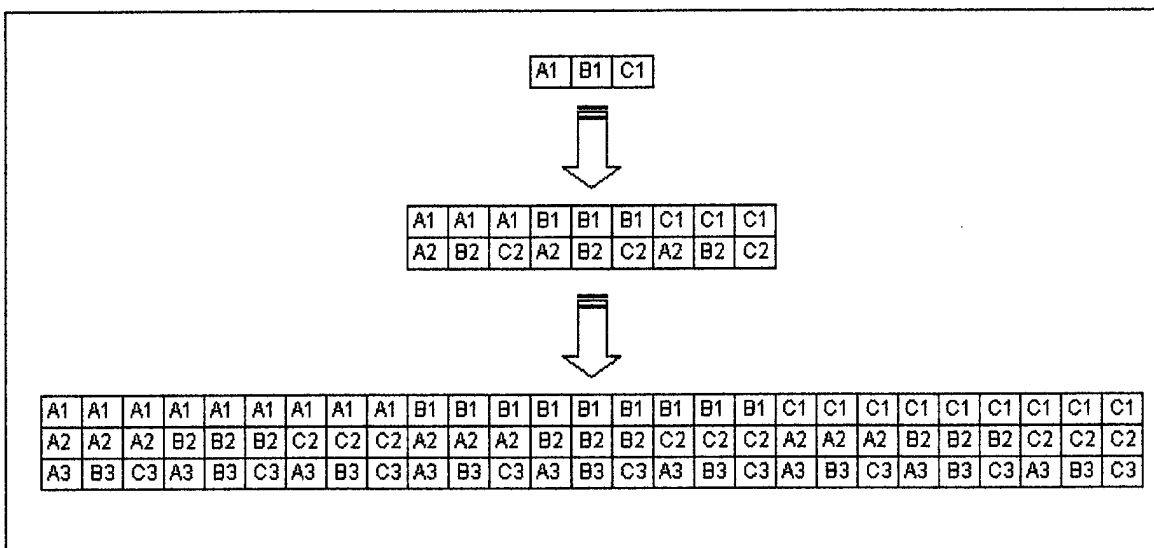
Fig. 11 - Generation of track table

A second table containing *M-1* rows is generated by taking the element by element differences between consecutive rows. The elements of this table represent changes in displacement from scan to scan for all possible tracks. The standard deviations of the elements in each column are computed. The columns of the table are sorted in the order of increasing standard deviation. Recall that each column corresponds to a track. The column with the smallest standard deviation corresponds to the track with the most constant speed. Tracks with the smallest standard deviations are considered to be the most probable tracks.

The algorithm was developed and tested using simulated data. Illustrated in Figure 12 are the positions, at uniformly spaced scan times, of a target traveling along a road at constant speed. Illustrated in Figure 13 are the scan by scan positions of five targets, each moving with constant (but different) speed over a road. Shown in Figure 14 by lines connecting the target positions from scan to scan, are the five most probable tracks selected by the algorithm. In the absence of "jitter" the five most probable tracks selected by the algorithm were indeed the true tracks. The presence of jitter in the data, however, can cause the algorithm to assign a higher probability to some false tracks than to some of the true tracks. This is illustrated in Figure 15 which shows the five most probable tracks for data with 30 meters of rms jitter. For this case one of the five most probable tracks is a false track. The fifth true track was the seventh most probable track found by the algorithm. The point at which jitter causes the selection of false tracks depends on the number, separation, and speed of the targets, as well as on the scan update rate. A clear indication of false tracks is the association of a single detection into more than one track. However, in the absence of collateral information, it is not possible to identify which of the tracks sharing a detection is a false track.

The algorithm was designed to be robust with respect to data dropout and unresolved targets. If target α is detected in all but one scan, and if the tracks encompass all scans, it is unavoidable that detections of target α will be associated with detections of some other target in that "bad" scan, forming a false track. Depending on the positions of the other detections in the bad scan, it is possible that all such false tracks for target α would have large standard deviations. If this is the case, the algorithm would be

Fig. 12 - Target traveling at constant speed along road

unable to track target α. A solution to the problem would be to skip over bad scans when forming tracks. The algorithm permits the user the option of skipping over specified scans when forming tracks. This option greatly enhances the robustness of the algorithm with respect to data dropouts and unresolved targets. Shown in Figure 16 are the results for a scenario consisting of five targets over five scans with one or more data dropouts in the second scan. The Figure shows the seven most probable tracks (all five real tracks plus two false tracks) selected by the algorithm by skipping scan 2. The inclusion of scan 2 would have resulted in the selection of several additional false tracks by the algorithm.


Fig. 13 - Five Targets Viewed Over Five Scans

Fig. 14 - Five most probable tracks identified by algorithm

Fig. 15 - Effect of position jitter on track association algorithm

Fig. 16 - Seven most probable tracks with second scan deleted

SUMMARY AND CONCLUSIONS

Two algorithms, a target location algorithm and a track association algorithm, were developed for possible use on the E-2C/APS-145. The target location algorithm uses DFAD road data to leverage the good downrange accuracy of the radar to substantially improve target location accuracy. The algorithm was evaluated using simulated data. The algorithm is susceptible to large errors if the separation between roads is small compared to the bearing accuracy of the radar. Such errors could be avoided if appropriate collateral information is available. The algorithm takes into account road segment orientation with respect to the LOS in the determination of MLE target position.

The track association algorithm takes advantage of the fact that the target location algorithm converts a two dimensional problem into a one-dimensional problem. The algorithm forms and tests all possible associations of target detections from scan to scan. The algorithm searches for tracks that correspond most closely to targets moving with constant speed. The algorithm was tested using simulated data and found to perform reasonably well in the presence of jitter in the position data. The approach used for dealing with data dropout and unresolved targets is to skip over selected scans when forming tracks.